ALCF Early Science Program

Argonne NATIONAL LABORATORY

ESP Kick-Off Workshop
Project Plan Presentation

Multiscale Molecular Simulations at the Petascale

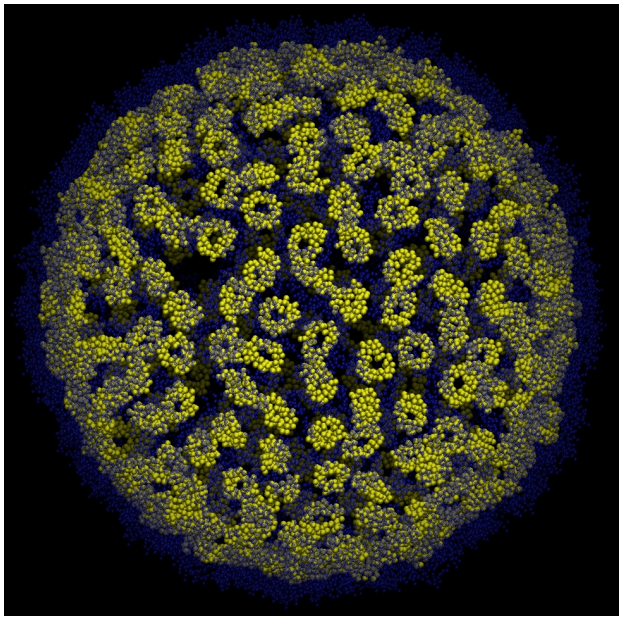

PI: Gregory Voth (U. Chicago)

Presenter: Christopher Knight

October 18-19, 2010

Project Overview

Cellular-scale molecular modeling of biological processes.

- **Multiscale approach will enable new bio-simulation of cellular-scale processes**
  - For the coarse-grained interactions studied, experimental data insufficient
  - Must be refined via proposed atomistic simulations that Mira enables

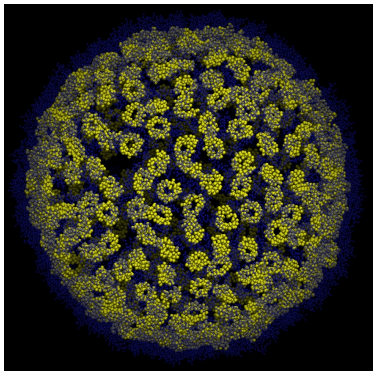

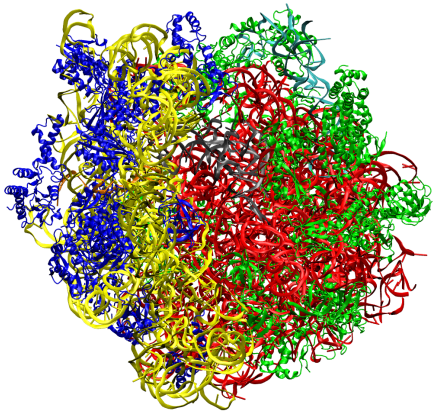

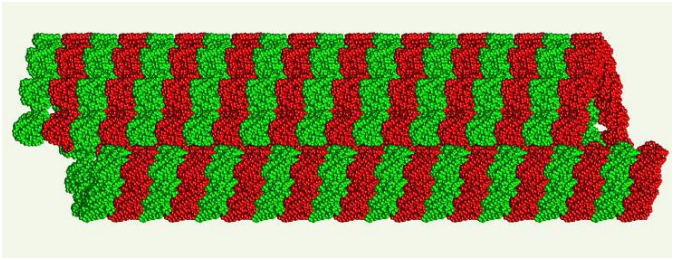

Scientific Field: Molecular/Chemistry

Codes: NAMD, LAMMPS. GROMACS, MSCGFM, TANTALUS

Computational Approach, Numerical Methods

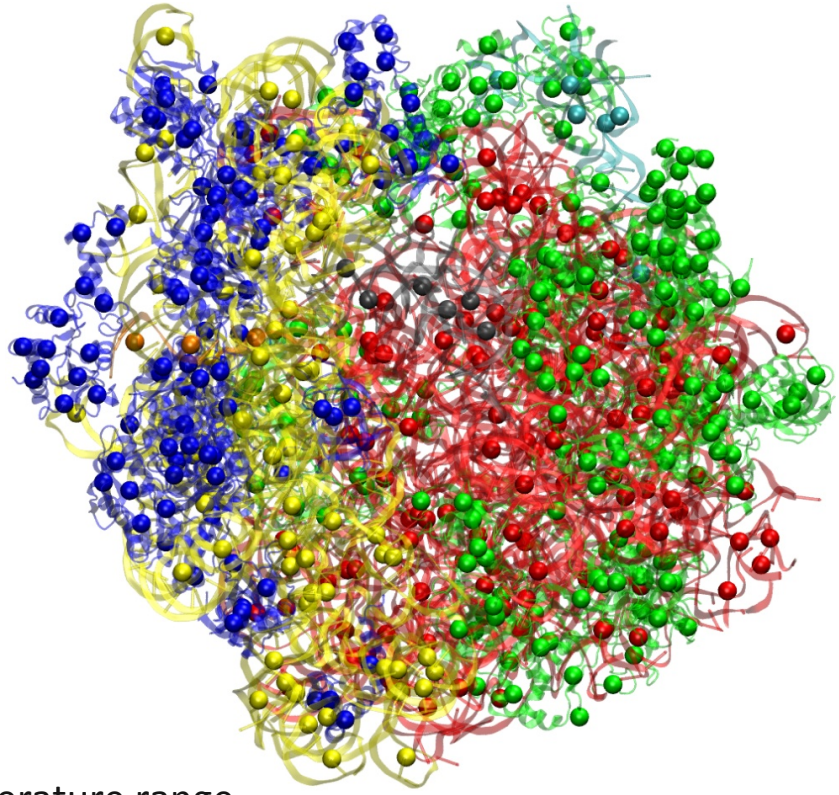

- **Multi-scale approach:**
  - Atomistic
  - Coarse-grained (CG)
- **Molecular Dynamics**
  - Solve Newton's equations of motion
  - Discrete time-steps
  - Nearest neighbor communications
- **Replica exchange**
  - Multiple 'copies' within pre-specified temperature range
- **Multiscale coarse graining (MSCG)**
  - CG model reproduces structural properties of reference all-atom data

Parallelism and Existing Implementation

- **TANTALUS**
  - spatial and linked-cell decompositions
- **MSCGFM**
  - Two-level parallelism
- **MPI everywhere (no threads)**
- **I/O: order-n per replica**
- **Current Performance/Scalability**
  - TANTALUS scales to 64,000 cores of NICS Kraken (Cray XT5) with 85% efficiency
    - 83.2 million CG particles
    - 1000 CG sites per core

Library and Tool Dependencies

- **Libraries**
  - MPI-2
  - FFTW
  - LAPACK
  - BLAS
  - ScaLAPACK
  - PETSc
- **Tools/other codes**
  - GROMACS
  - NAMD
  - LAMMPS

Anticipated Modifications for Blue Gene/Q

- **Two-layer parallelization of LS-solver in MSCGFM:**
  - Atomistic trajectories first subdivided into *blocks*
  - *Intra*- and *inter*-block parallelization levels
  - ScaLAPACK for dense intra-block LA (SVD)
  - PETSc for sparse intra-block LA (LSQR)
  - Tunable block sizes
  - Scalability is topology-dependent -> *per-case* optimization

- **Performance and scaling needed to run proposed problem on Mira:**
  - 0.1 sec/step on 70,000 cores

Plan for Next 6 Months Effort

- **Help find and hire a project postdoc**
- **Parallelization of MSCGFM**
  - Implementation via ScaLAPACK and PETSc
  - MPI-2 for inter- and intra-block communications
- **Tune block sizes for each CG case on Intrepid**
- **Use projections to estimate performance on Mira**
- **Investigate using CG representation to reduce I/O**
- **Explore hybrid MPI/OpenMP version of LAMMPS (ICMS)**